



Federal Aviation
Administration

Transforming Risk Management

**Understanding the Challenges
of Safety Risk Measurement**



The challenge of safety risk measurement is threefold:

To gather data pertaining to everything
that happens to a given flight

To understand, from departure to arrival,
the relationships among the factors
affecting that flight

To discern properly the activities
designed to ensure that flight's safety
from their actual outcomes

Contents

4	Introduction
6	Risk Analysis Processes (RAP)
8	Airborne RAP
11	Surface RAP
12	Service Integrity RAP
13	RAP Next Steps
14	Advances and Advantages in Digital Tools
15	Key Performance Indicators
17	Converging Runway Operations
19	Digital Tools Next Steps
21	Conclusion



: Introduction


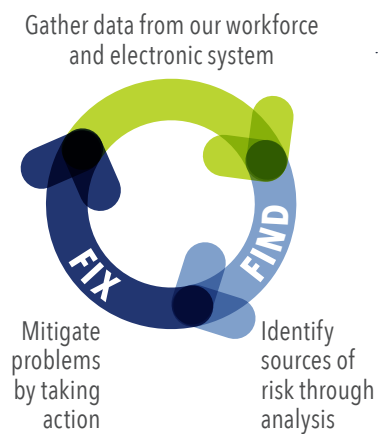


Figure 1. The Air Traffic Control Safety Performance Continuum

As the National Airspace System (NAS) becomes more complex and congested, the Air Traffic Organizations (ATO's) mission—to provide the safest, most efficient air traffic services possible—becomes correspondingly more challenging. Our ability to ensure the quality of our services requires a commitment to objectively measure, improve, and communicate safety performance.

Simply stated, the ATO must be very good at **collecting information and using that information to find and fix problems.**

Air traffic control safety performance measurement can be thought of as movement along a continuum (See Figure 1). It began with simple counts of procedural violations (e.g., losses of required separation)—an imprecise but entirely natural approach, given a data infrastructure consisting of little more than flight plans, time/speed calculations, position reports, and broadband radar.

The introduction of tracking systems and data recording capabilities provides more precise measurements and a broader pool of information to support investigations of accidents and their causes. This level of detail allows us to better understand—and, thus, focus our resources and actions on—how to prevent the development of major safety incidents and minimize the number of things that go wrong.




Figure 2. Advances in aviation safety from 1929–2015 include radar innovations and NextGen technology

Finally, at the other end of the continuum, advanced measurement techniques are adapted and expanded to embrace not only adverse events, but the complex dynamics of everyday operations. The purpose of these measurements is to more thoroughly understand how, under varying circumstances, procedures are properly executed and by what mechanisms hazardous situations are prevented from developing. This final phase helps us ensure that as many things as possible go right.

It is important to note that none of the phases in this continuum supersede the others; we must:

- 1** **Know** quickly when and where things are going wrong
- 2** **Understand** the causal factors and chains of events that lead to hazardous situations
- 3** **Apprehend** the full range of mitigations available to us, including those with the flexibility to support safe, efficient operations

Achieving these goals requires that we move beyond traditional risk identification and measurement techniques to leverage data generated by the new technologies that increasingly compose our infrastructure (e.g., digital communication, navigation, surveillance, and decision-support systems). By taking more subtle measurements of the NAS, we will advance our understanding of the factors that



contribute to or mitigate hazards, driving a more adaptive and anticipatory approach to safety.


This paper describes recent improvements to the ATO's safety performance measurement infrastructure. It focuses, specifically, on our enhanced Risk Analysis Processes and on the deployment of an entirely new set of safety intelligence tools—our Key Safety Performance Indicators.

: Risk Analysis Processes

The ATO's Risk Analysis Processes (RAPs) are designed to systematize the way that we analyze the most serious safety incidents in the NAS. RAPs employ panels of independent

subject matter experts who, supported by software-based tools, analyze and score the risk of these incidents. RAP tools standardize and quantify the panels' analyses.

There are currently **three versions of RAP:**



In their broad outlines, the RAPs address basic elements of aviation safety, elements that together represent an incident's risk—that is, the combination of an incident's

severity and likelihood of recurrence (or repeatability). For both Airborne and Surface RAPs, severity ratings are determined by calculating the proximity and closure rate of the involved vehicles (their risk of collision) and analyzing the performance of the NAS's defensive layers (barriers), which represent the controllability of the situation. Service Integrity RAP uses selections from a weighted look-up table to establish severity. Repeatability for all three versions of RAP is determined by analyzing any equipment, procedural, or management deficiencies that contributed to the incident (systemic factors), any human errors involved (non-systemic factors), and the likelihood that all of those factors will align again at some point in the future (the window of opportunity) (See Figure 3).




Figure 3. How Does RAP Work?

Reviewing all data available and relevant to an incident, a RAP panel considers each of these elements and, using a software tool that weights the contribution of the involved factors, generates a numeric risk score reflecting severity and repeatability. This numerical representation allows events to be situated in the context of other, similar events, forming datasets suitable for trend and comparative analysis. RAP datasets allow us to develop a more fully integrated picture of pilot, controller, and technician performance and more precisely targeted risk mitigation strategies; these datasets are also shared globally with other air navigation service providers.

Always striving to improve our safety performance, the ATO has continued to examine, improve, and expand our RAPs.

/// Airborne RAP

Incidents that occur *in flight*

+ The ATO has conducted several evaluations of **Airborne RAP**.

The results of these evaluations have revealed a number of issues, including: subjectivity and inconsistencies in the process's results, inaccuracies in its weighting scheme, and insufficiencies in its data attributes (i.e., the conditions and

causal factors available for selection by RAP panels).

To address these issues, Airborne RAP underwent a series of modifications. We expanded Airborne RAP's list of causal factors, developed a new set of scoring

: Airborne RAP

weights derived from NAS historical data, and, to address inconsistencies in RAP results, replaced the process's original barrier structure with a more refined barrier model. While the old structure produced only a single aggregate risk score, the new model allows us to see and measure the performance of the NAS's many layers of defense.

These barriers fall into three categories: air traffic control, pilot, and NAS technology infrastructure. Each category is composed of many discrete barriers designed to prevent a loss of required separation from occurring (termed resolution barriers) or prevent a loss of separation from becoming a collision (recovery barriers) (See Figure 4).

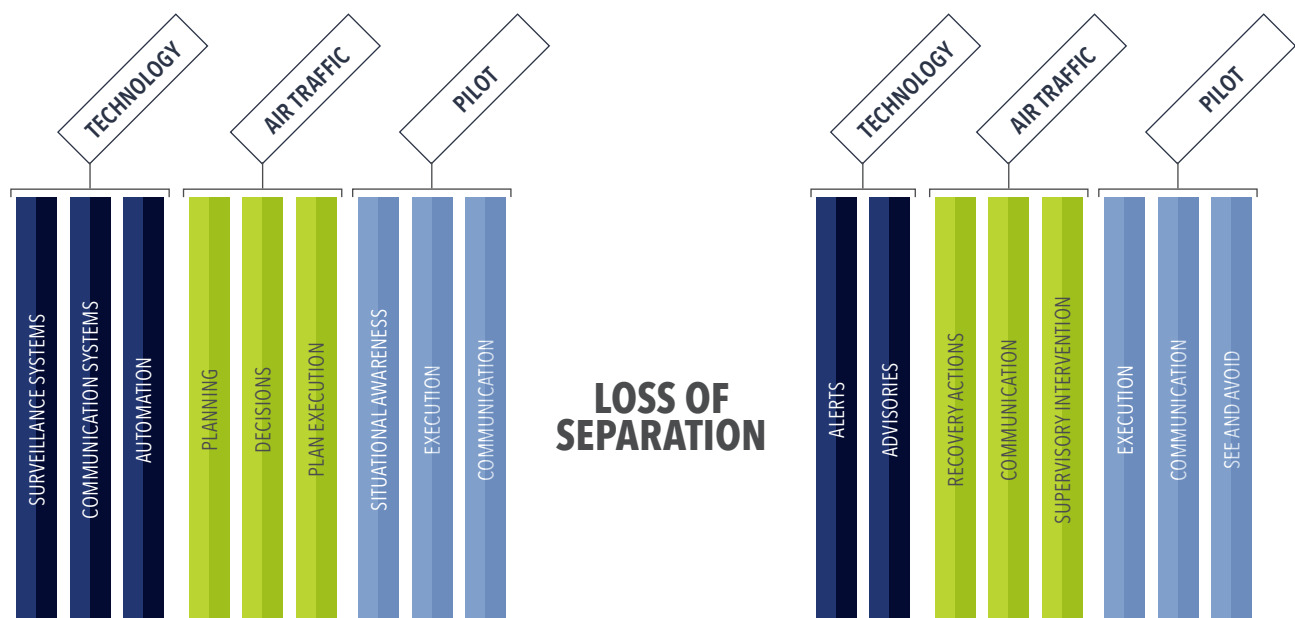


Figure 4. Airborne RAP's Barrier Model with Examples of Individual Barriers




Figure 5. Barrier Failures

Barrier failures represent degradations of the NAS's protections against actual collisions (See Figure 5). The failure

of any barrier increases the controllability score, and thus the severity score, of the incident under analysis.

The restructured Airborne RAP produces scores for each discrete barrier (e.g., conflict detection, controller plan, communication systems). These new scores allow safety analysts to inspect the effectiveness of barriers and the factors that influence their performance

at different locations and at different levels of specificity (ranging from airports, to regions, to the NAS as a whole). By grouping and analyzing these scores, we can identify trends in individual or composite barrier performance, which, in turn, help us understand the conditions that contribute to vulnerabilities or successes.

Adopting a barrier model has also improved the scalability of RAP data. Because risk scores are, in part, determined by the ratio of available to failed barriers, the scores do not depend on the stability of the NAS's current structure; the addition or removal of barriers at any point in the future will not affect the internal consistency of RAP data.

/// Surface RAP

Incidents that occur *on the runway*

: Surface RAP




Figure 6. Surface RAP Severity Structure, Showing Additional Severity Inputs and Flexible Barriers

/// Continuous Equations

Since weather is constantly changing, continuous equations are used to generate weather scores


+ In 2014, the ATO deployed Surface RAP, which is designed to analyze incidents that occur on airport movement areas. Like the restructured Airborne RAP, the ATO developed Surface RAP as an explicit barrier model. There are, however, several

important differences between the airborne and surface environments and, thus, between their respective RAPs. For example, while the airborne environment is composed only of fixed controllability barriers (e.g., pilot communication, controller decisions, system alerts, and other factors that are considered part of the NAS no matter where an incident occurs), the surface environment includes flexible physical and technological

barriers that may or may not be present at a given airport (for example, a specific runway lighting system).

Surface RAP also incorporates information pertaining to runway entry location, which is indicative of the energy state of the involved aircraft and the time available for controller or pilot reaction, and weather, including runway visual range (RVR) data (the distance a pilot can see down a runway) and ceiling data (the height of the cloud base from the ground)

+ (See Figure 6). **Continuous equations** are used to generate



weather scores for any combination of RVR and ceiling, allowing for more precise representations of the context in which the incident occurred.

To determine the most accurate way of distributing the possible scores for aircraft-to-aircraft proximity and closure rate, the ATO conducted statistical analysis of historical runway incursion data to **derive and map percentiles for both parameters to a 0-to-10 scale.** Because events involving vehicles and pedestrians constitute a very small portion of the current dataset, temporary scales, based on the available historical data, were established; these scales will be modified as more data becomes available.

Repeatability analysis in Surface RAP also differs from Airborne RAP because it is influenced by airport geometry and temporary surface conditions, such as runway closures and construction activities.

Completing the ATO's suite of RAPs is Service Integrity RAP, which, like Surface RAP, was rolled out in 2014. The

: Service Integrity RAP

purpose of this version of RAP is to assess the risk associated with any unexpected failure, interruption, or degradation

of NAS equipment or services that could affect the ATO's ability to provide safe air traffic control or flight information

/// Service Integrity RAP

Technology-related incidents

+ services. Service Integrity RAP promises to systematize the way that we identify technical issues that impact operations, allow us to better understand those issues, and guide our resources appropriately.

Although Service Integrity RAP follows the same basic principles as the other RAPs, the numerous and dynamic factors affecting the environment in which our technicians work mean that its severity subsection cannot rely on a barrier model. Instead, the severity of each incident is determined by the selection of failure combinations from a look-up table designed to represent all possible effects of system failure modes on operational safety. Service Integrity RAP's look-up table currently comprises more than 5,600 failure combinations.

All factors relevant to the incident under analysis are captured in a customized technical-occurrence taxonomy and examined to determine whether they are causal/contributory, observed (unusual, but of neutral effect), or positive.

: RAP Next Steps

Fine-tuning the RAPs is an ongoing process of perfecting the barrier models and look-up table, producing increasingly accurate scoring weights, and bolstering repeatability calculations with historical data. The Surface and Service Integrity RAPs are newly operational, and the

restructured Airborne RAP is currently in the final phases of testing.

: Advances and Advantages in Digital Tools

The ATO's investments in safety intelligence are producing state-of-the-art data recording, modeling, analysis, and visualization tools that we are using to monitor and measure NAS safety performance to a degree never before possible. By mining big data, we are able to discern, at a remarkable level of detail, the factors affecting air navigation safety and the relationship of those factors to emerging issues; this, in turn, enables us to more precisely focus our risk mitigation efforts and resources. Among the most important of our safety intelligence tools are Key Performance Indicators and tactical decision-support tools for air traffic controllers.



DIGITAL TOOLS USED TO MONITOR AND MEASURE NAS SAFETY PERFORMANCE

/// Key Performance Indicators

Algorithmic tools developed in partnership with The MITRE Corporation; also known as KPIs



Developed in partnership with The MITRE Corporation, **Key Performance Indicators (KPIs)** are algorithmic tools. Fed with threaded track data (which merge input from automation systems, operational safety databases, and flight tracking systems), they are used to monitor targeted flight parameters, operations, and environmental factors throughout the NAS.

Our KPI portfolio includes 20 algorithms, 14 of which have been fully deployed (See Figure 7). Their outputs populate

: Key Performance Indicators

two kinds of safety dashboards: management and analytical. Management dashboards provide near real-time overviews of safety

trends in the NAS, including the frequency and location of specific types of operations and incidents. Analytical dashboards enable inquiry into the details of specific incidents. A subset of these dashboards is used by local safety councils, who are often in the best position to make safety decisions for their facility.

The following case study, *Converging Runway Operations*, illustrates how the ATO is using cutting-edge technologies to identify issues and improve safety performance.

The High Energy Approaches KPI uses various surveillance and airport configuration data to calculate the specific energy profile of aircraft on final approach. It then compares that profile to a distribution of flights reflecting aircraft type, airport environmental factors, weather factors, and runway length, flagging as "high energy" any profile that exceeds the nominal limits by more than two standard deviations. All data relevant to the flagged events are saved to a database for analysis (See Figure 8).




Figure 8. High-Energy Approach Profile



KPI	Algorithm	Dashboard
Missed Approaches	●	●
High Energy Approaches	●	●
Final Approach Overshoots	●	●
Class-B Airspace Excursions	●	●
Risk of Runway Overrun	●	●
Rejected Takeoffs	●	●
Simultaneous Parallel Approaches: Overshoots and No Transgression Zone Blunders	●	●
Arrival Winds Overview	●	●
Similar Call Signs	●	●
Airport Daily Overview	●	●
En Route Daily Overview	●	●
Terminal Radar Approach Control Daily Overview	●	●
Terminal Traffic Alert and Collision Avoidance System (TCAS) Resolution Advisory (RA) Simulator	●	●
Departure Procedure Compliance	●	
En Route TCAS-RA Simulator	●	○
Converging Runway Operations	●	○
Altitude Deviations	○	○
Opposite Direction Operations	●	●
Minimum Vectoring Altitude Violations	○	○
Top 5 Hazards	●	○

Initial Development Begun ○
Prototyped ●
Completed ●

Figure 7. Status of the ATO's Current KPIs

: Converging Runway Operations




Figure 9. Converging Runway Operations

/// Monte Carlo Simulation

A broad class of computational algorithms that rely on repeated random sampling to obtain numerical results; they are often used in physical and mathematical problems and are most useful when it is difficult or impossible to use other mathematical methods

During a 2011 investigation of go-arounds – i.e., situations in which an aircraft on final approach aborts its landing and

climbs away from the runway – the ATO discovered an issue related to Converging Runway Operations (CRO): the potential for the flight path of an aircraft executing a go-around to conflict with the flight path of a departing aircraft (See Figure 9). An unrelated analysis effort conducted in 2012 revealed a number of CRO events with elevated risk levels. In response, the ATO issued a Corrective Action Request with recommendations for procedural

and technological improvements. Three more incidents occurred between January and July 2013, and, in July 2013, the National Transportation Safety Board issued a Safety Recommendation highlighting five CRO incidents at three airports in which safety was unacceptably degraded.

The ATO undertook a number of initiatives to address the problem. Beginning in 2011, we partnered with The MITRE Corporation to conduct in-depth analysis of the issue. Using four years of surveillance and aeronautical data, we developed a set of representative rejected-landing trajectories and then put that set into a **Monte Carlo simulation**. The resulting model helped us to identify the specific scenarios that posed the most risk.

The ADW is a graphical box, displayed on tower controllers' monitors, that demarcates an area on the final approach path to a runway. Its purpose is to indicate to controllers when it is safe to release aircraft for departure—i.e., only when there are no arrival aircraft within the ADW (See Figure 10).




Figure 10. The Arrival Departure Window

We then re-purposed the model in two ways:

- 1** **Developed** a CRO KPI (now available on our dashboards)
- 2** **Created** a site-specific decision-support tool for air traffic controllers, known as the Arrival-Departure Window (ADW) (See Figure 10)

Two additional steps, taken to address procedural concerns related to CRO, began in 2013:

- 3** **Required** all air traffic managers at airports with CRO geometry to convene safety risk management panels to review their operations
- 4** **Formed a joint FAA-industry workgroup** tasked with prioritizing those airports and developing an implementation plan for new safety requirements which included:
 - A policy dictating that CRO be conducted dependently (meaning that controllers must coordinate departures from and arrivals to runways with intersecting flight paths)
 - Procedures for alternating runway configurations
 - The adoption of both revised configurations and the ADW at 24 high-priority airports, which was adapted and optimized for each site using runway configuration and NAS operational data

: Digital Tools Next Steps

The success of the ATO's safety intelligence program is continuing with new tools and KPIs. One of our most recent tools is the [REDACTED]

/// CROPD

Speech-recognition technology to decipher voice commands; it generates an alert if the controller issues a clearance for a take-off or landing on a closed runway

Developed in response to a recent increase in runway incursions on closed runways, CROPD uses speech-recognition technology to decipher controller clearances, accepts controller input on which runways are closed, and generates an alert if the controller issues a clearance for a take-off or landing on a closed runway.

Among several new KPIs, the Opposite Direction Operations (ODO) KPI has shown significant promise. It was developed in response to unsatisfactory efforts to adequately mitigate the risks posed by situations in which landing operations are conducted on the same runway as, but in the opposite direction of, take-off operations.

Figure 11 represents one of the ATO's ODO monitoring dashboards. In the ODO map view, users can easily see the relative frequency of ODO across the NAS; hovering over a site on the map reveals airport-specific information.





Figure 11. Map Depicting Relative ODO Rates

The ODO KPI utilizes an algorithm designed to detect and measure occurrences of aircraft with converging flight paths (head-on within 45 degrees) when departing and arriving on the same runway



The System Risk Event Rate (SRER) Trendline from October 2011 to June 2015

Operational safety risk is measured by the SRER, which has been trending down over the past four years as a result of proactive safety initiatives.

while within 10 nautical miles of each other laterally, and less than 1,000 feet vertically.

These are just a few of the ways we are benefiting from our investments in safety intelligence. Over the course of the coming year, we will likely add to our set of KPIs, developing alongside them the dashboards and decision support tools that will allow us to improve the safety of the services we deliver.

Figure 12 represents an analytical dashboard depicting the status of ODO events by runway. Each square represents a runway at a specific airport; the size of the squares indicates



Figure 12. ODO Heat Map Dashboard

the number of ODO events and the color, the most severe event at that runway (determined by aircraft proximity). Each square on the dashboard is labeled with the airport identifier, runway name, and number of ODO incidents.

An analyst can select an airport from the heat map to bring up additional views with event-specific information, including aircraft identification, departure/arrival runways, time-of-event, vertical/lateral separation, altitude profiles, ground speed, and aircraft track position.



Figure 13. ODO Aircraft Track Position

Another analyst view is depicted in Figure 13, one of many available for examining the details of an ODO event. The view provides graphical information on arriving and departing aircraft tracks; the closest proximity off the aircraft is depicted in red.



: Conclusion

While refining and expanding the ATO's approaches to safety measurement has been challenging—requiring significant time and resources to develop, evaluate, and optimize new analytical processes and tools—our commitment to keeping

ahead of the latest developments in air traffic congestion and

safety technology integration is paying off. As a result of the trends and insights that the RAPs and KPIs provide, we are deepening our knowledge of the factors and sequences of events that contribute to or, alternatively, mitigate safety hazards in the NAS. Armed with this new knowledge, we are developing increasingly adaptive and anticipatory approaches to safety, and deploying ever more targeted and effective risk mitigation strategies.



Joseph Teixeira

Vice President, Safety and Technical Training

Air Traffic Organization, Federal Aviation Administration





**U.S. Department of Transportation
Federal Aviation Administration**

Federal Aviation Administration

Air Traffic Organization
Safety and Technical Training
800 Independence Avenue SW
Washington, DC 20591

